

Riding the Waves of Supercomputing Technology

Looking toward the future, Livermore's Computation Directorate leverages three high-performance, cost-effective technology curves to deliver computing capability at a faster rate than described by Moore's Law.

BUILDING supercomputers powerful enough to run complex simulations of the performance of nuclear weapons is a challenge that the National Nuclear Security Administration's (NNSA's) Advanced Simulation and Computing (ASC, formerly ASCI) Program has met full on since its beginning in 1995. But even as the capabilities of the ASCI machines—Red, Blue, White, Q, and the upcoming Purple—have grown, so have the sophistication of the applications and the number of users eager to run their simulations.

Livermore's users, like their counterparts at Los Alamos and Sandia, include not only weapons scientists who run simulations for NNSA's Stockpile Stewardship Program, but also researchers whose simulations of, say, plasma instabilities or the performance of insensitive high explosives feed vital information into stockpile stewardship efforts. In addition, other scientists working on research important to Livermore's overall national security mission—for example, nuclear nonproliferation, detection of underground structures, nondestructive evaluation, earthquake hazard analysis,

or oil exploration—want to run simulations on ASCI-caliber supercomputers available through Livermore's Multiprogrammatic and Institutional Computing (M&IC) Initiative. With ASC and ASCI-class machines, there always seems to be a need for more—more capability to run scientific calculations at large scale and more capacity to process a varied workload from many users simultaneously.

As Michel McCoy, head of the Integrated Computing and Communications Department (ICCD), notes, Livermore is in many ways a victim of its own success. "Because three-dimensional code development funded by ASC and other programs has been so successful at Livermore," he explains, "the demand for supercomputing capability and capacity to explore complex scientific issues has become enormous. Someday, even the 100-teraflops [trillion floating point operations per second] Purple machine will not be powerful enough to help us answer these big science questions. The demands of classified stockpile stewardship simulations, which have

the highest priority, can crowd out the classified basic science calculations related to stockpile stewardship that run on ASCI machines. We're also faced with the issue of how to accommodate unclassified basic science applications that are, nevertheless, important to the ASC Program's mission. Add to all of this the science problems from a multitude of Laboratory programs that can't be solved without help from ASCI-class supercomputers, and the need for more supercomputing capability and capacity increases further."

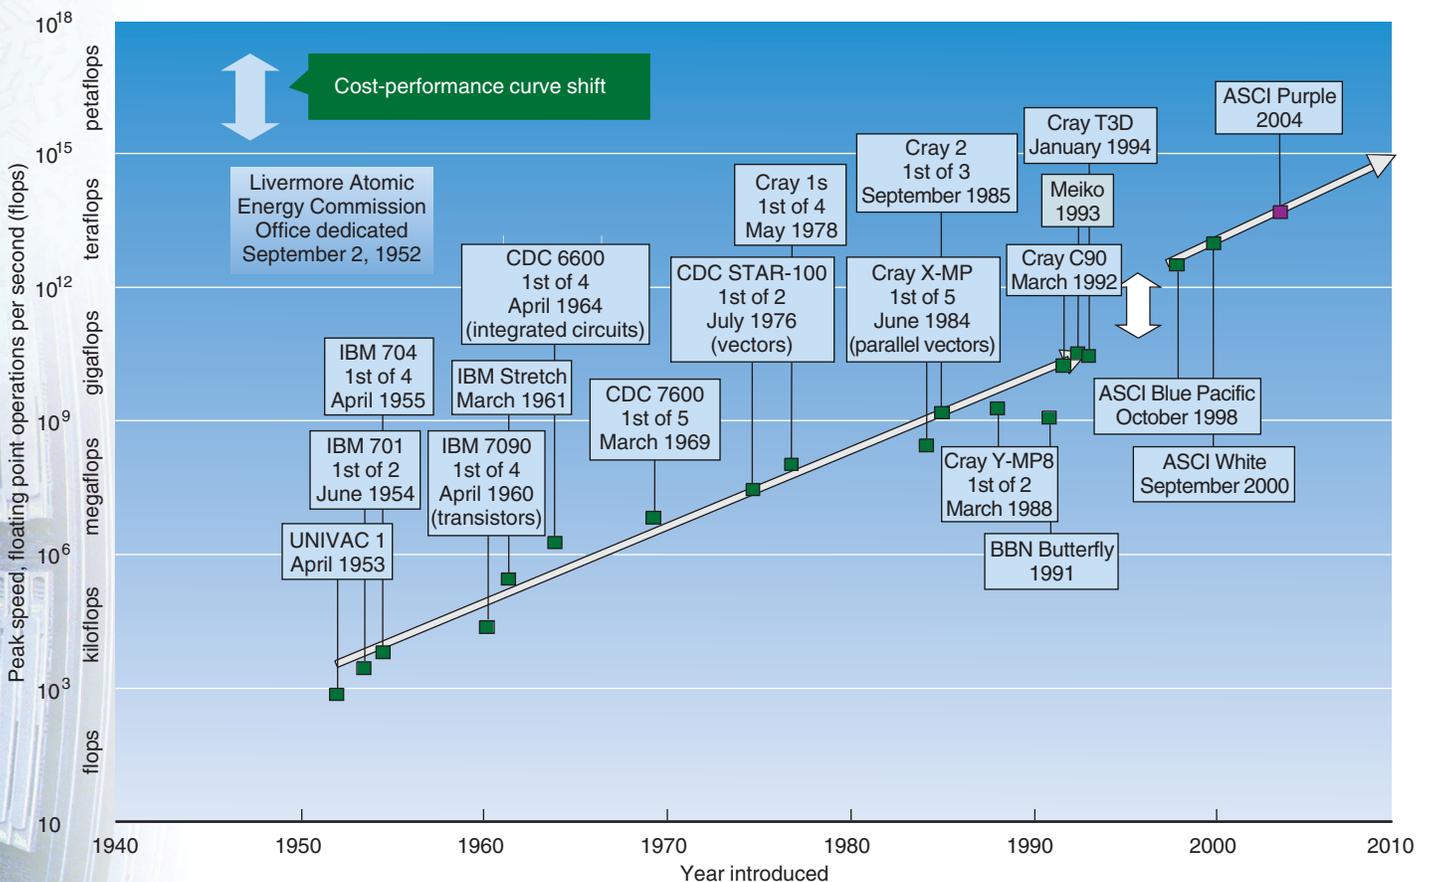
The only way to generate the type of capability required, says McCoy, is to deliver computing capability faster than the rate described by Moore's Law, which states that processing power

doubles every 18 months. And the rate must be faster not just by a little, but by a lot. "If we continue to follow the existing technology curve, namely, vendor-integrated multiprocessor platforms, we will not get to multiple petaflops by 2010."

To meet the ever-growing needs of the ASC Program and the rest of the scientific community, Livermore's Computation Directorate is implementing a computing strategy that promotes switching to and straddling new cost-performance computer technology curves, or waves—a balancing act of timing, prescience, and prediction. To get where it needs to be and deliver capacity and capability at low cost, Livermore must jump from

the wave of present technology to the next new technology wave at the right time—and ride not just one wave, or two, but three of these technology waves simultaneously. Each wave can provide benefits to some area of scientific research, even when the technology is new and unproven. As the technology matures, the new system becomes useful to additional kinds of research.

McCoy points out that the tri-laboratory ASC Program and the institutionally funded M&IC program have cooperated and leveraged expertise to exploit the various capability-capacity technology curves. In particular, ASC has funded Livermore's Computer Center and its



The history of supercomputing at Livermore includes jumps between technology curves to gain cost effectiveness and increased speed and capability. If supercomputing continues on the present curve, it will approach a quadrillion floating point operations per second (petaflops) by 2010 but will not reach the goal of multiple petaflops.

large base of expertise in the field of ASCI computers, and the Laboratory has leveraged this expertise and funded alternative high-performance, cost-effective supercomputing technologies for institutional users. "This synergy benefits both parties," says McCoy. For ASC, it accelerates the rate at which the new technologies mature. For the institution, it provides cost-effective supercomputer capacity and capability across the Laboratory.

This type of synergy is not new at Livermore. Earlier, the Laboratory worked with Compaq and Quadrics to develop the TeraCluster2000 (TC2K) which is part of M&IC. (See *S&TR*, October 2001, pp. 4–12.) The partnership with Compaq on TC2K made possible Compaq's successful bid for ASCI Q at Los Alamos and the computer funded by the National Science Foundation at the Pittsburgh Supercomputing Center.

Riding the Technology Waves

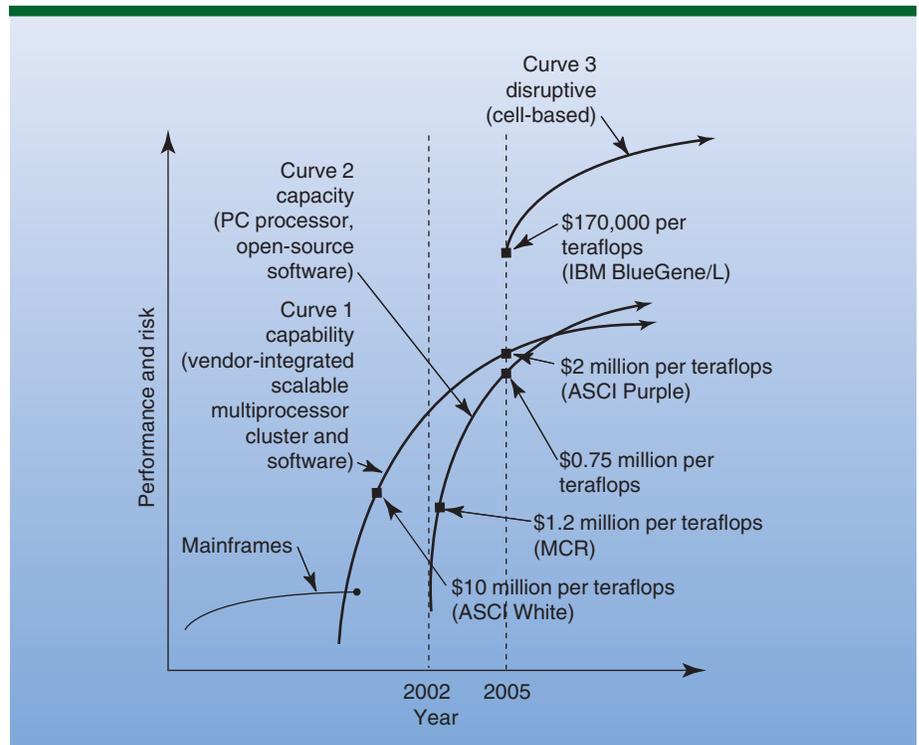
Livermore has made jumps from one computing technology to another in the past. First were the mainframes, followed in the 1970s by vector supercomputers and in the late 1980s by massively parallel processing supercomputers. With the cessation of underground testing of nuclear devices in 1992 and the birth of the Stockpile Stewardship Program the following year came the need for much better computer simulations to help ensure that the nation's nuclear weapons stockpile remained safe, reliable, and operational.

The ASC Program was created to provide the integrating simulation and modeling capabilities and technologies needed to combine new and old experimental data, past nuclear test data, and past design and engineering experience. The result is a powerful tool for future design assessment and

certification of nuclear weapons and their components. ASC required machines that could cost-effectively run simulations at trillions of floating point operations per second. This requirement forced a sea change in the supercomputing industry and a jump to another technology wave—massively parallel scalable supercomputers.

The ASCI machines use many thousands of reduced instruction set computer (RISC) processors—a class of processors found in workstations—working in unison instead of the more expensive, one-of-a-kind specialized processors characteristic of earlier parallel processing. ASCI machines delivered "more bang for the buck." Code developers benefited as well. "In about six years," adds McCoy, "we went from systems where one-dimensional codes were routine and two-dimensional codes were possible, but a stretch, to ASCI systems where

Since any given computing technology curve is ultimately limited by Moore's Law, Livermore's Computation Directorate is embracing a strategy to straddle and, when the time is right, switch to new technology curves. Vendor-integrated massively parallel ASCI machines are the current workhorses of the ASC Program. The Linux cluster machine typified by Multiprogrammatic Capability Resource (MCR) will lead to the next-generation production systems. Cell-based technology, such as that to be used in BlueGene/L, appears to be an affordable path to petaflops systems. Only time will tell whether other cost-effective paths will emerge that might lead to the petaflops regime.



2D is the norm and 3D calculations are done, but are sometimes a stretch.”

By all accounts, the ASC Program has been remarkably successful in meeting its goals. Its machines are the workhorse computers for running the complex two-dimensional and three-dimensional codes used to meet the nation's most demanding stockpile stewardship requirements. “These machines and their codes are proven,” says Mark Seager, assistant department head for platforms in ICCD. “They are used routinely for stockpile stewardship and for must-have deliverables. These machines and the codes that run on them are reliable and trusted, both of which are a necessity for stockpile stewardship. There's no room for error when simulating nuclear weapons.”

With each ASCI system, the cost per teraflops fell as technology advanced. For Livermore's ASCI White, the cost was \$10 million per teraflops. For ASCI Q, a Los Alamos machine, the cost is \$7 million per teraflops. For ASCI Purple, an upcoming Livermore machine, the estimate is \$2 million per teraflops. But now, this particular technology curve—vendor-integrated systems using high-performance workstation processors and proprietary vendor software—has matured to the point where the cost per teraflops improves at only a slow exponential rate as dictated by Moore's Law. So people such as McCoy and Seager are looking at technologies that can be exploited at better cost performance after the 100-teraflops Purple. They are also interested in building smaller-capacity machines at much lower cost.

Two new technologies hold promise. A near-term technology based on cluster architecture—that is, large groups of interconnected commodity microprocessors (the type found in desktop personal computers and laptops)—combined with open-source,

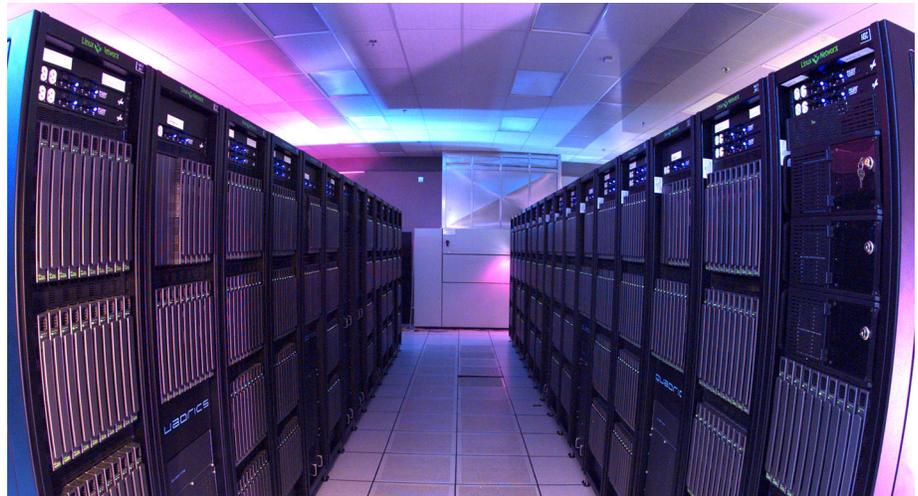
or nonproprietary, software is epitomized in the Multiprogrammatic Capability Resource (MCR) system. Farther in the future are cell-based supercomputers that would use system-on-a-chip technology and low-cost, low-power embedded microprocessors. This third technology is embodied in BlueGene/L, a system designed by IBM and slated for delivery to Livermore in December 2004.

Next on the ASCI Crest

ASCI Purple will fulfill the goal set in 1996 to achieve 100 teraflops by mid-decade for stockpile stewardship use. Purple will have a peak performance (100 teraflops) equivalent to 25,000 high-end personal computers. It will have 50 trillion bytes of memory and 2 petabytes of disk storage capacity, the equivalent of a billion books or about 30 times the contents of the Library of Congress. About the size of two basketball courts, Purple will have more than 12,000 IBM Power5 microprocessors. It will allow scientists to run a full-physics, full-system model of a nuclear weapon in three dimensions.

ASCI Purple will be eight times more powerful than ASCI White, which was the first ASCI system powerful enough to investigate crack propagation in the materials of a nuclear weapon in three dimensions. A major step forward, Purple will allow designers and code developers to focus increasingly on improving physics models. For instance, if weapon designers need to change some element of the weapon, they can insert this change into the initial conditions of the simulation, and the recalculation will show them the effect of the changes. One three-dimensional simulation representing the operation of the weapon system for a fraction of a second will still take about eight weeks of computing time.

Sometime between October and December 2003, the Early Delivery Technology Vehicle (EDTV) will be available at Livermore as part of the Purple contract. EDTV will consist of at least 32 nodes and feature the new IBM Federation switch, the node interconnect that is being used on Purple. “Having EDTV here first will help us avoid some issues we faced



The Multiprogrammatic Capability Resource at Livermore combines open-source software and cluster architecture to provide Advanced Simulation and Computing-level supercomputing power for unclassified research.

Simulating Laser-Plasma Interactions

Using the Laboratory's recently installed Multiprogrammatic Capability Resource (MCR) system, physicists Steven Langer and Bert Still spent the early part of 2003 simulating an experiment scheduled to occur this summer on the National Ignition Facility's (NIF's) Early Light (NEL) system. NEL consists of the first four beams of NIF and was first fired at high power in December 2002.

The experiment simulated by Langer and Still will study laser-plasma interactions under conditions similar to those that will be found in hohlraum targets in future NIF ignition experiments. (A hohlraum is a small, cylindrical chamber used to enclose the target in experiments on high-power lasers.) A NIF laser beam has to cross roughly 5 millimeters of plasma between the laser entrance hole and the gold wall of the hohlraum. The laser heats the wall so much that the wall emits x rays in a configuration designed to uniformly bombard a spherical capsule at the center of the hohlraum, squeezing it down until it reaches fusion temperatures and pressures.

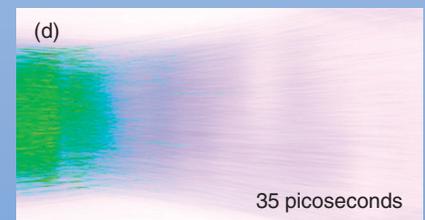
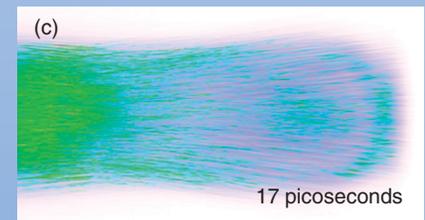
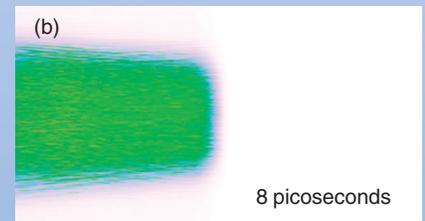
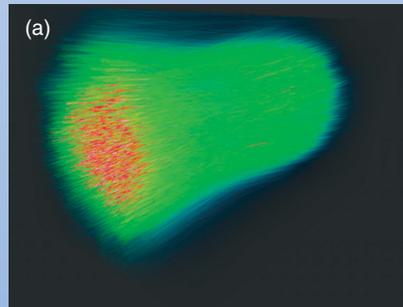
The experiment will use a 4.5- by 4.5- by 5.5-millimeter plastic bag filled with neopentane gas. The gas inside the bag will be quickly ionized, creating a plasma that will be a reasonable surrogate of that in the actual NIF target, and the NEL beams will

pass through it. It is this interaction of light and plasma that Still and Langer sought to simulate with the code PF3D, which Still and others began developing in the mid-1990s.

In the simulation, the laser beam passes through the central portion of a 1- by 1- by 5-millimeter plasma. "Our physics algorithms constrain a computational cell, or zone, to be only slightly larger than the wavelength of the laser light, which is 0.35 micrometers," explained Langer. "We ended up using 6.8 billion cells in the calculation, which is an enormous number of cells."

The simulation ran for 10 days on 1,920 of MCR's processors to simulate 35 picoseconds, the time it will take for the laser light to travel 5 millimeters twice. Results from the simulation occupy about 14 terabytes of archival storage, which, Langer notes, was roughly 25 percent of all the information written to storage at the Laboratory in February 2003. None of Livermore's existing volume visualization programs could handle a 6.8-billion-zone simulation with a lot of fine-scale structure. Mark Duchaineau, supported by the Advanced Simulation and Computing Program's Visual Interactive Environment for Weapons Simulation project, wrote a software volume visualization program that would handle the data.

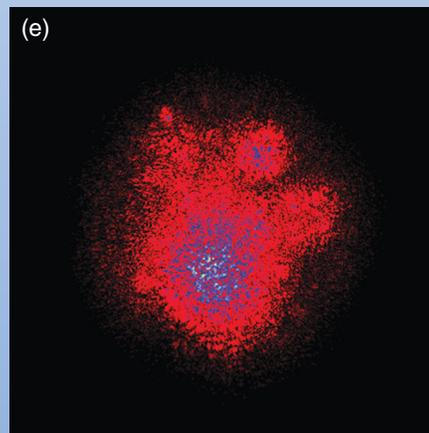
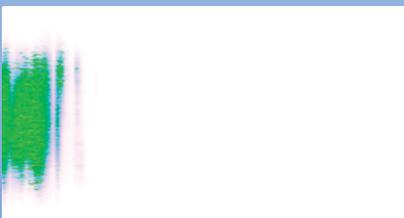
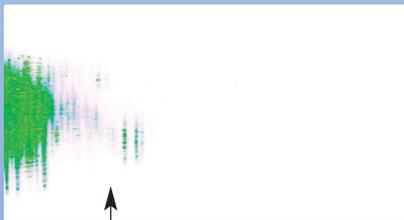
This volume visualization of a simulation done on the Multiprogrammatic Capability Resource shows a National Ignition Facility laser beam as it passes through a hot plasma. A cross section of the laser beam has many small, bright spots. These bright spots form long, thin filaments as they pass through the plasma. The lowest intensities are blue, intermediate intensities are green and yellow, and the highest intensities are red. (a) This oblique view shows that the brightest (red) filaments are concentrated in the core of the beam. (b) This side-on view shows that after 8 picoseconds, the laser (left) has penetrated halfway through the plasma. Scattering (right) is weak. (c) After 17 picoseconds, the laser beam (left) has almost reached the back of the simulation volume. Scattered light (right) is strong, and most of the laser light is absorbed or scattered (as intended in the experimental design). (d) After 35 picoseconds, the transmitted laser light (left) is spread into a wider cone, and the scattered light (right) is still strong. (e) Scattered light comes in bursts. The arrow in the scattered light (right) part of (c) points at one burst. The view in (e) shows the intensity of the burst in a plane perpendicular to the beam direction. Low intensity in (e) is shown in red, intermediate intensity in blue, and high intensity in white. The scattered light starts as a narrow spot and becomes wider as it passes through the plasma. The largest blob in (e) becomes nearly as wide as the simulation volume by the time it leaves the plasma volume.



The results, notes Langer, were a real eye-opener. "We found that the laser beam penetrates most of the way through the 5 millimeters of plasma, then the backscattered light becomes strong, and the forward beam broadens as it transverses the plasma. We can see in detail how the bursts of scattered light grow stronger as they move through the plasma."

On the basis of their simulation, Langer and Still can make some predictions about the upcoming experiment that will help designers fine-tune NEL systems. For example, they can provide information about where to put detectors on the target chamber walls to pick up transmitted light and what detector sensitivity settings are appropriate for measuring backscattered and transmitted light.

"This was a great example of how, at Livermore, simulation and experiment work hand-in-hand," says Langer. "The work of planning and building NIF and developing ASCI computers and massively parallel computer codes such as PF3D are now coming to fruition. We have the results of simulations, showing us what we expect will occur in the gas bag experiments and providing information that will help guide design of the first NEL experiments. We're eager to see what the data show when the experiments run this summer."



with White, when switch, node, and file systems were all brand new to us at one time," says McCoy. Delivery of all of Purple's 197 refrigerator-size processing units is scheduled to be completed by December 2004. Once up and running, Purple will be the primary supercomputer for the tri-laboratory ASC Program and a production resource to stockpile stewardship.

Catching the Next Wave

The time to move to the next technology wave, one that can deliver increasingly cost-effective capability and capacity for running the next generation of simulations, is nearing for the ASC Program.

Currently, ASCI machines use vendor-designed and -maintained operating software, computers, and processors. All that service doesn't come cheap. "We pay for this service and intellect," says McCoy. So the second technology curve, which Livermore is now exploring at scale using institutional funding, features open-source, not vendor-proprietary, software; cluster architecture; and microprocessors of the kind found in personal computers and laptops (32-bit Pentium-4 Xeon processors, in particular, although 64-bit processors are also being evaluated).

"The result is that we won't have to buy the software, and we'll be using much less expensive components, but we'll own all the problems," notes McCoy. "Owning problems is expensive and psychologically unsettling, of course, so there are pluses and minuses to this approach. But our experience, as we move forward cautiously, is that it is better to be in control than to be dependent. If this fails, we will also know whom to blame."

The first step is to build capacity systems and mid-level capability systems with open-source cluster

technology. Livermore is securely positioned on this second computer technology wave with its new 11.2-teraflops MCR system. It nearly matches the 12.3-teraflops ASCI White in power, but at \$1.2 million per teraflops, its cost per teraflops is 10 times less than ASCI White's.

MCR is a 32-bit microprocessor-based cluster built by Linux NetworX for M&IC. Cluster computers are composed of many identical or similar types of machines that are tightly coupled by high-speed networking equipment and message-passing interface software. The MCR cluster uses an open-source software environment based on Linux and the Lustre global file system.

"Using open-source software," says Seager, "means we can often fix or address our own problems rather than hand them over to a vendor, which we have to do for systems that run proprietary software. The vendor,

however, may not even be able to reproduce our problems because of a difference in scale between the vendor's systems and ours. So MCR is a much more efficient proposition for us." Seager adds that the other advantage to running open-source software is that many open-source developers are creating codes and software that end up in a big pool of open-source development. Jumping into that pool means that finding areas of mutual interest and setting up collaborations may become easier.

MCR arrived at Livermore in the summer of 2002 and started running large-scale applications in December 2002. MCR is now running select unclassified science simulations for stockpile stewardship and other Livermore programs. The system will go into full production mode as soon as the new file system is stabilized. McCoy points out that MCR's arrival represented a second example of the

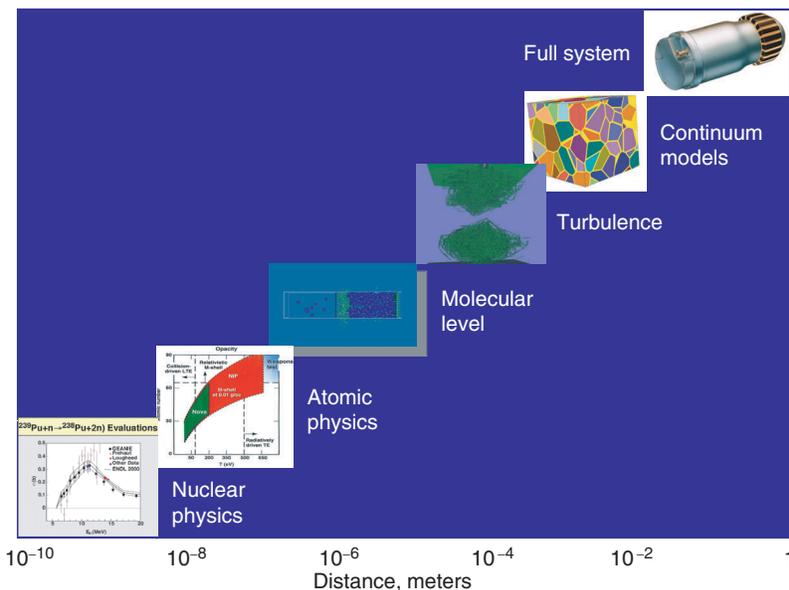
synergy between the Laboratory and the ASC Program. "The MCR system incorporates a new technology ideal for basic science investigations, but it is not yet ready to approach the scale of problems presented by ASC weapons codes," he says. "Yet, MCR would not have been possible without ASC's investments in technology and the expertise developed at Livermore."

The goal is to provide scientists throughout Livermore with stable Linux clusters for a general scientific workload. "MCR is a capability and capacity machine for running large-scale parallel scientific simulations. The plan is to get it into the hands of scientists quickly," adds Seager.

The system was upgraded to its current size in early 2003. Since that time, a handful of scientific teams have been using it to run codes, large and small, to help get the system ready for general availability at 11.2 teraflops this summer. Among the early users were physicists Bert Still and Steven Langer, who used MCR to simulate one of the first experiments scheduled to be performed on the National Ignition Facility this summer. (See the [box on p. 8.](#))

Geophysicist-computer scientist Shawn Larsen also acquired time on MCR to run his seismic wave analysis code E3D for exploring issues related to test readiness, nuclear nonproliferation, and earthquake prediction. (See the [box on p. 11.](#))

"This open-source, cluster curve will almost certainly provide the transition to the next generation of ASC computers," says McCoy. "We plan to ride this curve now and shake down these new, unproven machines with our unclassified science codes. Then, in the coming months when we're convinced it's ready, we can shift this technology to stockpile stewardship by introducing moderate-size clusters into the classified environment for mid-sized problems."



BlueGene/L may offer a rapid path to the multiple petaflops level, allowing scientists to finally address a number of complex physics issues from the atomic scale (on the order of nanometers) to the macroscale (on the order of meters).

New System Heightens Three-Dimensional Reality in Seismic Simulations

Geophysicist-computer scientist Shawn Larsen was a member of one of the science teams that put the Multiprogrammatic Capability Resource (MCR) through its paces when it first became available in December 2002. He used his E3D code, a powerful seismic code that incorporates three-dimensional information about the propagation of seismic waves, to get a more detailed picture of how these waves interact with different geologies and topographies in their path. His results contribute to a number of Livermore's efforts, including test readiness, earthquake hazard analysis, and nuclear nonproliferation.

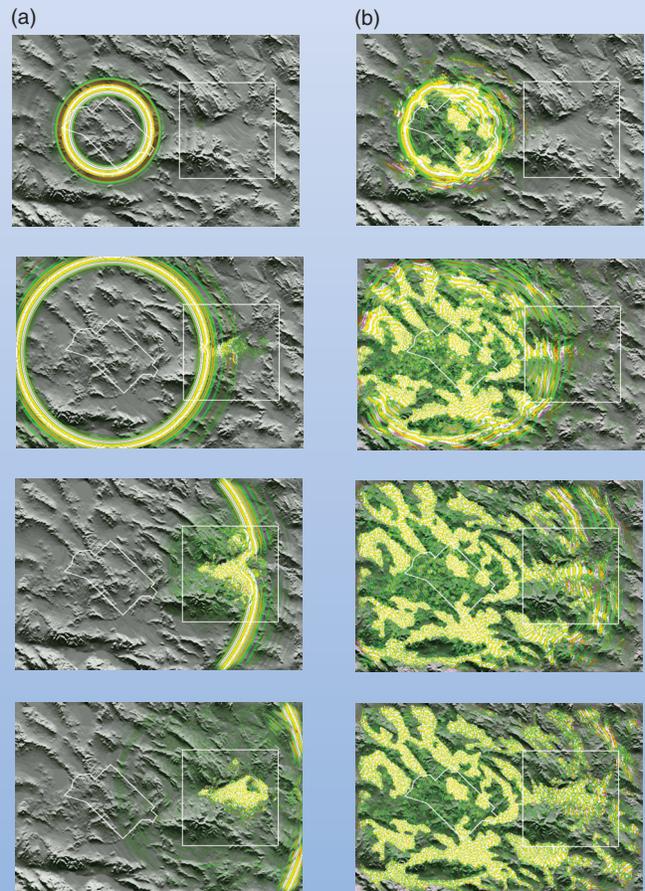
A 1993 Presidential Decision Directive requires that the national laboratories shall be ready within a specified amount of time if the nation decides to resume underground nuclear testing. Larsen's task was to use a three-dimensional geologic model developed at the University of Nevada at Reno to simulate the seismic shaking that would occur in the Las Vegas area from a nuclear underground test at the Nevada Test Site. "Las Vegas sits in a basin," says Larsen. "Since the last test in 1992, Las Vegas has grown considerably to the north, into deeper parts of the basin. We wanted to look at the different types of seismic waves and see how they all propagated from the source of an explosion, through the intervening geology, to the basin."

Running simulations that include hills, mountains, valleys, and other topography required a larger supercomputer than previously available. With the arrival of MCR, this kind of complex three-dimensional simulation became possible. (See [top figure at right.](#)) "For practical purposes, most of these simulations need over 6 billion zones," explains Larsen. To complete the 16 simulations needed required about 300 hours and 1,600 processors. "We could have easily used the entire machine," says Larsen, "but we left some of it free for others." Each simulation used about 1.5 terabytes of memory.

Using MCR, the team observed details that had never before been revealed by modeling. The energy from the point source of the simulated underground explosion no longer radiated in neat rings, as in previous calculations. It scattered. Also, the topography caused some of that radial energy to "leak" into the transverse direction, creating ground motion perpendicular to the radial direction. "This is the first time we've seen this leaking in a model," says Larsen. "We have seen it in the data and measurements, and now we're finally able to simulate it. The topographic simulations indicated that ground motions can be amplified at the tops of ridges, hills, and mountains, which is important to know when looking at earthquake hazards." (See [bottom figure at right.](#))

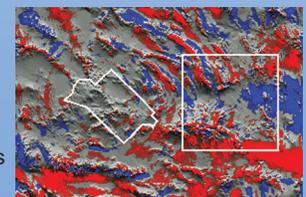
For the nuclear nonproliferation effort, knowing how seismic energy propagates is also important. The emphasis is on understanding the seismic signal: Does it come from a mining blast

or underground nuclear explosion? "We need to understand the complexity of propagation, and how topography affects propagation," he said. "The kind of modeling we're doing on MCR brings us one step closer to being able to fully understand and discriminate a seismic source."



Seismic energy moving in (a) a simple geologic model and (b) a Multiprogrammatic Capability Resource model using complex geology. Note the vastly increased level of detail in (b).

Results from a Multiprogrammatic Capability Resource simulation show seismic shaking on the tops of hills and in valleys, as well as the more intense shaking on ridges compared with that in valleys.



Growing on the Horizon

The third wave, cell-based computer technology, is even farther out on the technical horizon. It shows great promise of yielding machines that are even more powerful and less expensive than their predecessors. However, since the technology is unproven, investing in this technology today involves high risk. For example, cell-based technology would feature systems using low-cost, low-power embedded microprocessors. Embedded microprocessors, Seager notes, are everywhere—in cars, CD and DVD players, telephones, and

other consumer electronics. Livermore and IBM are working together to use this technology to produce systems on a chip—that is, to combine most of the features of a node (microprocessor, memory controller, network, for instance) into a single chip, or cell, instead of the many chips that perform these functions in present-day ASCI computers. Combining node features in a cell reduces power consumption and cost and permits more nodes to fit in a given amount of floor space.

This third curve will be realized through BlueGene/L, a computational

sciences research and evaluation machine that IBM will build in parallel with ASCI Purple and deliver in 2005. BlueGene/L will be used for unclassified research into areas such as first-principles molecular dynamics for materials science, three-dimensional dislocation dynamics of materials, high-explosives, and turbulence.

Like MCR, BlueGene/L will run the open-source Linux operating system to take advantage of all the pluses in the open-source arena of code development. The machine will be based on

Supercomputer Characteristics						
	Curve 1 technology			Curve 2 technology	Curve 3 technology	Currently world's fastest
Machine	ASCI White	Q	Purple	MCR	BlueGene/L	Earth Simulator
Machine peak speed, teraflops	12.3	20.0	100.0	11.1	180.0 to 360.0	40
Total memory, terabytes	8	33	50	4.6	16 to 32	10
Footprint, square meters	930	1,266	1,115	465	230	3,200
Total power, megawatts	1.0	2.5	4.5	0.3	1.2	10.0
Cost, millions of dollars	100	160	240	13.9	less than 100	about 350 (Does not include maintenance or nonrecurring engineering costs.)
Number of nodes	512	2,730	197	1,152	65,536	640
Total central processing units	8,192	10,922	12,608	2,304	130,000 microprocessors	5,120

Key characteristics of the computers belonging to Livermore's three technology curves compared with each other and with those same features of the Earth Simulator, currently the world's fastest computer.

130,000 advanced microprocessors and have a theoretical peak computational rate of 367 teraflops, with a cost per teraflops of \$170,000. When it's up and running, it promises a major advantage in peak speed over present-day computers, including Japan's Earth Simulator, currently the fastest supercomputer. BlueGene/L could be the "next bigthing." If all goes well, the next-generation system, BlueGene/P, could be the first petaflops machine performing a quadrillion calculations per second.

Scientists and researchers are looking forward to BlueGene/L and pondering what it may mean to them and their research. Larsen, who does seismic wave modeling, says, "On MCR, we can create a limited number of complex simulations, one after the other. BlueGene/L will allow us to do hundreds or even thousands of simultaneous simulations, in which we can more easily vary the physical parameters by considering many types of geologies using a statistical approach."

Physicists Langer and Still are also intrigued with the possibilities of BlueGene/L. "By industry standards, the average life of supercomputers like MCR is about 200 weeks," says Langer. "So simulations like our 10-day run on laser-plasma interactions for the National Ignition Facility can't monopolize the machine, much as we'd like to. Others are waiting their turn. But BlueGene/L has the potential to make runs such as ours routine by the time NIF comes fully online in 2008."

Balancing on Technology Waves

"We see cell-based BlueGene/L delivering an affordable means to petaflops supercomputing, which is where we need to be in 2010," says Seager. "But we have to stay open to

other possibilities and not commit entirely too early. In the next few years, other technology curves may become apparent—disruptive technologies we can't predict—that will lead to a breakthrough regime by the year 2006 or 2007."

Disruptive technologies are those that change the game quickly and unexpectedly—the Internet, for example. The trick is to gauge when the time is right—if it ever is—to switch to a disruptive technology. "Some technologies make it, some don't, and it's important not to switch too early to something that may not be there in the longer term," says Seager. However, says McCoy, "An institution such as the Laboratory needs to pursue at least one of these new approaches whether or not it works out in the end, because experimentation is necessary for evolution."

For the ASC Program and for the Laboratory, the goal is to cost-effectively deploy advanced, high-performance computing architectures. The first wave of these systems is the current reliable one, with the massively parallel, scalable ASCI Purple next in line.

The second wave, which is based on open-source codes and cluster architecture, is embodied in Livermore's MCR. In January 2003, MCR was ranked the fifth fastest supercomputer in the world on the TOP500 supercomputing list, the first time ever that a computer based on Linux cluster technology broke into the list's top 10. In addition, it was the only Linux-based supercomputer to appear in the top 5, leading industry experts to note that MCR is an important step in supercomputing history because it demonstrates the potentially large effect Linux clusters will have in the high-performance computing community.

The emergence of Linux clusters is just the latest wave on the horizon, gaining speed as it roars to prominence. As high-performance technical users such as Lawrence Livermore and researchers such as Still, Langer, and Larsen move to clustered solutions, the technology will be tested and it will mature in reliability and functionality.

The third wave, which includes cell-based design using system-on-a-chip technology and embedded microprocessors, will be explored in BlueGene/L. With petaflops potential, this machine promises to open a new universe of scientific simulation. Seager says, "Having Blue Gene/L will be like having an electron microscope when everyone else has a magnifying glass."

When Purple and BlueGene/L are both fully operational, which is expected before the end of 2005, they will, combined, have a higher theoretical peak capacity than the 500 fastest supercomputers in the world today. Then, it will be time to catch that breaking wave and ride it to shore.

—Ann Parker

Key Words: Advanced Simulation and Computing (ASC) Program, ASCI Purple, ASCI supercomputers, BlueGene/L, computation strategy, E3D, laser-plasma interactions, Linux clusters, Multiprogrammatic Capability Resource (MCR), National Ignition Facility Early Light (NEL), PF3D code, seismic wave analysis.

For further information contact Michel McCoy (925) 422-4021 (mccoy2@llnl.gov).